

BEHAVIORAL COPYSerial No.: 10/024,553
Docket No.: 0023-0062Remarks

Applicant appreciates the Examiner's indication that claims 8 and 15 are directed to allowable subject matter. Further, in the Office Action, the Examiner objected to the disclosure based on an informality; rejected claims 1-5 and 12 under 35 U.S.C. § 102(e) based on U.S. Patent No. 6,781,992 to Rana et al. ("Rana"); rejected claims 18-24 and 26 under 35 U.S.C. § 102(e) based on U.S. Patent No. 6,832,261 to Westbrook et al. ("Westbrook"); rejected claims 6, 7, 9-11, 13, 14, 16, and 17 under 35 U.S.C. § 103(a) as being unpatentable over Rana in view of U.S. Patent No. 6,026,477 to Kyker et al. ("Kyker"); and rejected claim 25 under 35 U.S.C. § 103(a) as being unpatentable over Westbrook in view of Kyker.

By this Amendment, Applicant has amended the specification to correct the error noted by the Examiner. The objection to the specification is therefore traversed.

Rejection of Claims 1-5 and 12

Claims 1-5 and 12 stand rejected under 35 U.S.C. § 102(e) based on Rana. Applicant respectfully traverses this rejection.

Claim 1 is directed to a reordering device that includes a reorder buffer, a reorder buffer pointer, and a reorder engine. The reorder buffer is configured to store information relating to data items, each of the data items being associated a sequence number chosen from a sequence number space, the sequence number indicating a relative order of the data items. The reorder buffer pointer is configured to store a value indicating a rearmost active entry of the reorder buffer relative to the

Serial No.: 10/024,553
Docket No.: 0023-0062

sequence number space. The reorder engine is configured to receive the data items and classifying each of the received data items based on the sequence number of the data item and the value of the reorder buffer pointer. The reorder engine determines whether to store the information relating to the received data items in the reorder buffer based on the classification.

A proper rejection under 35 U.S.C. § 102 requires that a single reference teach every aspect of the claimed invention either explicitly or impliedly. Any feature not directly taught must be inherently present. See M.P.E.P. § 2131. Applicant submits that Rana does not disclose or suggest each of the features recited in claim 1.

Rana, for instance, does not disclose or suggest the reorder engine recited in claim 1. More specifically, Rana does not disclose or suggest, for example, a reorder engine configured to classify each of received data items based on sequence number of the data item and a value of the reorder buffer pointer, as recited in claim 1.

As described in the pending specification, a reorder engine, by classifying data items as recited in claim 1, can efficiently reorder streams that occasionally include highly out-of-order data items. (See Specification, at, for example, paragraphs 10, 23, and 37-55).

Rana describes a queue engine for reassembling and reordering data packets in a network. (Rana, Abstract). Rana discloses an IP reordering unit 34 that reorders out of sequence PDUs. (Rana, col. 5, lines 20-30). Neither the IP reordering unit 34 of Rana, nor any other element in Rana, however, discloses or

Serial No.: 10/024,553

Docket No.: 0023-0062

suggests a reorder engine configured to classify each of received data items based on a sequence number of the data item and a value of the reorder buffer pointer, as recited in claim 1. The reorder engine of claim 1 also determines whether to store the information relating to the received data items in the reorder buffer based on the classification. Rana also does not disclose or suggest these features of claim 1.

In rejecting claim 1, the Examiner contends that Rana discloses classifying received data items as recited in claim 1. Specifically, the Examiner states:

Fig. 1, element 44 (IRU memory) is associated with a reorder engine. IP reordering unit 34 uses IRU memory 44 to keep track of windows which reflect PDUs belonging to a particular traffic flow and where each PDU belongs in sequence. Keeping track of windows belonging to a particular traffic flow is associated with the reorder engine classifying each of the received data items based on the sequence number of the data item and the value of the reorder buffer pointer, the reorder engine determining whether to store the information relating to the received data items in the reorder buffer based on the classification. See col 5, lines 22-24.

(Office Action, page 3). Applicant respectfully disagrees with the Examiner's interpretation of Rana. Using a memory 44 to keep track of windows that reflect PDUs belonging to a particular traffic flow, as disclosed by Rana, in no way discloses or suggests the reorder engine recited in claim 1. If anything, Rana appears to disclose that PDUs are assigned to different windows based on their assigned traffic flow, not based on a sequence number of the data item and a value of the reorder buffer pointer, as recited in claim 1. Accordingly, Rana can not be said to disclose or suggest a reorder engine configured to classify each of received data items, as recited in claim 1.

Rana, at column 5, lines 53-64, discloses additional details of receiving data units and writing them to memory:

Serial No.: 10/024,553
Docket No.: 0023-0062

The process then proceeds to blocks 52 and 54 where PDU assembler 26 extracts the required information from the various PDU headers, and the PDU is written into packet memory 24, respectively. As described above, PDU assembler 26 can use the extracted header information for a variety of purposes, including but not limited to determining a unique identifier for the session with which the PDU is associated, determining whether the PDU is out of order, and determining whether the PDU is a fragment. After PDU assembler 26 extracts the required fields from the headers of the PDU, block 56 determines whether PDU assembler 26 has received the entire PDU. If the entire PDU has not been received then the process returns to block 52.

(Rana, col. 5, lines 53-64). In this section, Rana discloses using extracted PDU header information for a variety of purposes, including determining a unique identifier for the session and determining whether the PDU is out of order. Rana, however, appears to always write the PDUs into packet memory 24 after extracting the required information from the various PDU headers. Rana, however, does not classify PDUs in the manner recited in claim one, much less determine whether to store information relating to received data items in a reorder buffer based on the classification, as is also recited in claim 1.

For at least these reasons, Applicant submits that Rana does not disclose or suggest each of the features recited in claim 1. Accordingly, the rejection based on Rana is improper and should be withdrawn. The rejections of claims 2-5 based on Rana are also improper, at least by virtue of the dependency of these claims from claim 1, either directly or indirectly.

Independent claim 12 also stands rejected under 35 USC section 102(e) based on Rana. For at least the following reasons, Applicant respectfully traverses the rejection of claim 12.

Serial No.: 10/024,553
Docket No.: 0023-0062

Claim 12 is directed to a method of ordering data items comprising receiving the data items, the received data items being associated with a sequence number chosen from a sequence number space, the sequence number indicating a relative order of the data items. The method of claim 12 further includes classifying the received data items as one of a plurality of possible classifications based on the sequence number of the data items and a value that defines a position of a reorder buffer in the sequence number space. The method of claim 12 further includes processing the data items based on the classifications.

Rana does not disclose or suggest each of the features recited in claim 12. Rana, for example, does not disclose or suggest classifying the received data items as one of a plurality of possible classifications based on the sequence number of the data items and a value that defines a position of a reorder buffer in the sequence number space. As previously discussed with respect to claim 1, Rana does not classify data items based on a sequence number and a value that defines a position of a reorder buffer. If anything, Rana appears to disclose that PDUs are assigned to different windows based on their assigned traffic flow, not based on a sequence number of the data item and a value that defines a position of a reorder buffer. Rana, therefore, cannot be said to disclose or suggest each of the features recited in claim 12.

Accordingly, the rejection of claim 12 based on Rana is also improper and should be withdrawn.

Rejections of Claims 6, 7, 9-11, 13, 14, 16, and 17

Serial No.: 10/024,553

Docket No.: 0023-0062

Dependent claims 6, 7, 9-11, 13, 14, 16, and 17 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Rana in view of Kyker. Applicant respectfully traverses the rejections of these claims.

Dependent claim 6, for example, recites that the classifying utilizes a plurality of regions including a first region corresponding to the sequence number of the data item falling within a sequence number range covered by the reorder buffer, and a second region corresponding to the sequence number of the data item falling within a sequence number range extending from a head of the reorder buffer to a predetermined range beyond the head of the reorder buffer.

The Examiner, in rejecting claim 6, concedes that Rana does not disclose the features of this claim, but contends that Kyker discloses the features of claim 6, and that one of ordinary skill in the art would have found it obvious to modify the system of Rana by using the features as taught by Kyker. (Office Action, pages 10 and 11).

Kyker is directed to a branch recovery mechanism. (Kyker, Title). Applicant submits that Kyker in no way discloses or suggests classifying data items using the first and second regions that are explicitly recited in claim 6. The examiner appears to contend that Kyker discloses the first region at column 2, line 33 through column 3, line 21. (Office Action, page 11). The section of Kyker states:

FIG. 3 illustrates the operation of the instruction pipeline 10 of FIG. 1. Instructions are fetched, decoded into micro-ops, and stored in the ROB 18 to await execution. Steps A-C illustrate some examples of the type of operations performed in the instruction pipeline 10. At step A, a micro-op is stored at sequence number 4 in the ROB 18 (ROB entry 4) and the End-of-Buffer pointer 48 is incremented from ROB entry 4 to ROB entry 5 (where the next decoded micro-op will be stored in the ROB 18).

After the micro-op of sequence number 1 (ROB entry 1) has completed

Serial No.: 10/024,553
Docket No.: 0023-0062

execution, the execution result is stored in the ROB entry 1 and ROB entry 1 is marked as executed (a 1 is marked in the Executed bit for ROB entry 1). At step B, the micro-op of ROB entry 1 is retired by copying the micro-op's execution result from ROB entry 1 into the processor's real register set. The micro-op and other information in ROB entry 1 is then deleted and the start-of-buffer pointer 46 is incremented to point to ROB entry 2, which is the oldest (unretired) micro-op in ROB 18. ROB entry 1 is now available to receive a new micro-op.

At step C, the branch micro-op at sequence number 2 is executed and validated. In this example, during validation, it is determined that the branch of sequence number 2 has been incorrectly predicted, and therefore, the instructions that were prefetched and stored in the ROB after the branch entry 2 (micro-ops for ROB entries 3 and 4) are incorrect (were mispredicted). The micro-ops for ROB entries 3 and 4 correspond to the mispredicted path 75. In addition, the instructions in the pipeline stages earlier than the ROB 18 (e.g., the Instruction Fetch Units 12, the prefetch streaming buffer 40, the Decode stages 14, the RAT 16, the ID Queue 42) also correspond to the incorrect path and must be flushed. After flushing the micro-ops in the pipeline stages earlier than the ROB, the instructions corresponding to the correct path are then fetched and decoded. However, the ROB 18 cannot be flushed and the new micro-ops corresponding to the correct path cannot be loaded from the Front End section 30 into the ROB 18 until all instructions prior to and including the mispredicted branch (at ROB entry 2) have been executed and retired. When the mispredicted branch has been retired, then the ROB 18 is flushed or cleared and the micro-ops corresponding to the correct path can be loaded into the ROB 18.

In the example of FIG. 3, there are no unexecuted instructions before the mispredicted branch. Generally, however, because the branch operation may have been executed out of order (before one or more older instructions), the Front End section 30 can stall (wait) at the out of order boundary and cannot load the new (correct path) micro-ops from the Front End section 30 (e.g., Decode stages 14) into the ROB 18 until the mispredicted branch instruction has been retired and the ROB 18 flushed. (The out of order boundary refers to where the in order Front End section 30 meets the out of order Middle section 32). In other words, the Pentium does not mix correct and incorrect micro-ops in the ROB 18.

(Kyker, col. 2, line 33 through col. 3, line 22). This section of Kyker appears to relate to a branch prediction technique. Applicant fails to see how this section of Kyker in

Serial No.: 10/024,553
Docket No.: 0023-0062

any way discloses or suggests classifying data units using a plurality of regions, including the first region recited in claim 6. The examiner also cites this section of Kyker as disclosing the second region recited in claim 6. Applicant similarly fails to see how this section of Kyker in any way discloses or suggests classifying data units using a plurality of regions, including the second region recited in claim 6. This section of Kyker does not mention classifying data units or a reorder buffer, much less, as recited in claim 6, "a first region corresponding to the sequence number of the data item falling within a sequence number range covered by the reorder buffer, and a second region corresponding to the sequence number of the data item falling within a sequence number range extending from a head of the reorder buffer to a predetermined range beyond the head of the reorder buffer."

For at least this reason, Applicant submits that Kyker does not disclose or suggest the features of claim 6, and that Kyker therefore does not cure the deficiencies of the Rana with respect to claim 1. Accordingly, the rejection of claim 6 is also improper and should be withdrawn.

Claim 7 depends from claim 6 and further recites that the plurality of regions include a third region corresponding to the sequence number of the data item falling within a sequence number range outside of the first and second regions. The Examiner relies on Kyker to disclose the features of claim 7, and particularly points to element 34 of Kyker (Fig. 3) and column 1, lines 38-47. (Office Action, page 11).

Applicant respectfully disagrees with the Examiner's interpretation of Kyker.

Column 1, lines 38-47 of Kyker states:

The first seven stages of the instruction pipeline 10 (fetch 12, decode 14, RAT 14 and ROB 16 stages) are known as the In-Order Front End

Serial No.: 10/024,553
Docket No.: 0023-0062

section 30 of the processor because the instructions are kept in strict program order. The Dispatch stage 20 and the Execution stage 22 are known as the Middle Out-of-Order section 32 of the processor because micro-ops can be executed in any order. The Retirement stages 24 are also known as the In-Order Rear End section 34 of the processor because micro-ops are retired in program order.

(Kyker, col. 1, lines 38-47). This section of Kyker discusses an instruction pipeline. This section of Kyker in no way relates to classifying data units using a plurality of regions, including the third region recited in claim 7. If the Examiner maintains this rejection, Applicant requests that the Examiner clarify how the Examiner is interpreting Kyker to disclose the features of claim 7.

For at least this reason, Applicant submits that Kyker does not disclose or suggest the features of claim 7, and that Kyker therefore does not cure the above-discussed deficiencies of Rana. Accordingly, the rejection of claim 7 is also improper and should be withdrawn.

Applicant submits that the rejections of claims 9 and 10 based on Rana and Kyker are also improper, at least by virtue of the dependency of these claims from claims 6 and 7, respectfully. Further, the rejection of dependent claim 11 based on Rana and Kyker is also improper, at least by virtue of its dependency.

Dependent claims 13, 14, 16, and 17 also stand rejected based on Rana and Kyker.

Dependent claim 13 further defines the features of claim 12 and recites that classifying the data items further includes classifying the data items as belonging to a first region when the sequence number of the data items fall within a sequence number range encompassed by the reorder buffer, and classifying the data items as belonging to a second region when the sequence number of the data item falls

Serial No.: 10/024,553
Docket No.: 0023-0062

within a sequence number range extending from a head of the reorder buffer to a predetermined range beyond the head of the reorder buffer. Applicant submits that Kyker does not disclose or suggest these features of claim 13.

In rejecting claim 13, the Examiner points to the sections of Kyker quoted above in the discussion of claim 6 (i.e., column 2, line 33 through column 3, line 21). (Office Action, page 14). As discussed, this section of Kyker appears to relate to branch prediction techniques. In no way, however, do these sections of Kyker in any way disclose or suggest classifying data items as belonging to a first region or classifying data items as belonging to a second region based on the specific conditions recited in claim 13. If the Examiner maintains this rejection, Applicant requests that the Examiner clarify how the Examiner is interpreting Kyker to disclose the features of claim 13.

Claim 14 depends from claim 13 and recites that classifying the data items further includes classifying the data items as belonging to a third region when the sequence number of the data items falls within a sequence number range outside of the first and second regions. The Examiner relies on Kyker to disclose the features of claim 14 and points to the sections of Kyker quoted above in the discussion of claim 7 (i.e., column 1, lines 38-47). (Office Action, page 15). This section of Kyker in no way relates to classifying data items using a plurality of regions, including the third region recited in claim 14. If the Examiner maintains this rejection, Applicant requests that the Examiner clarify how the Examiner is interpreting Kyker to disclose the features of claim 14.

Serial No.: 10/024,553
Docket No.: 0023-0062

For at least these reasons, Applicant submits that Kyker does not cure the admitted deficiencies of Rana with regard to claims 13 or 14, and that therefore, the rejections of these claims are improper and should be withdrawn. The rejections of claims 16 and 17 are also improper, at least by virtue of the dependency of these claims from claims 13 and 14, respectively, and should also be withdrawn.

Rejection of Claims 18-24 and 26

Claims 18-24 and 26 stand rejected under 35 U.S.C. § 102(e) based on Westbrook. For at least the following reasons, Applicant respectfully traverses these rejections.

Independent claim 18 is directed to a network device comprising a data transmission component and a plurality of processing elements connected by the data transmission component, the processing elements communicating with one another by transmitting data items over the data transmission component. The processing elements each includes a reorder component configured to arrange received data items into an order corresponding to an order in which the data items were transmitted. Each of the reorder components includes a reorder buffer configured to store the data items, each of the data items including a sequence number chosen from a sequence number space, the sequence number indicating the order of the data items. Each of the reorder components also includes a reorder engine configured to classify the data items based on the sequence number of the data items and based on a position of the reorder buffer relative to the sequence

Serial No.: 10/024,553
Docket No.: 0023-0062

number space, the reorder engine determining whether to store the data items in the reorder buffer based on the classification of the data items.

Westbrook is directed to distributed resequencing and reassembly of subdivided packets. (Westbrook, Title). Westbrook discloses the distributed reassembly of large packets split into smaller packets with each packet marked with a sequence number, timestamp, or other order and reassembly indications. (Westbrook, Abstract). Westbrook, however, does not disclose or suggest each of the features recited in claim 18.

Westbrook, for example, does not disclose or suggest the reorder components recited in claim 18, which are each recited as including a reorder buffer and a reorder engine, where the reorder engine is configured to classify the data items based on the sequence number of the data items and based on a position of the reorder buffer relative to the sequence number space, the reorder engine determining whether to store the data items in the reorder buffer based on the classification of the data items.

In rejecting claim 18, the Examiner states that Westbrook discloses that the reorder engine recited in claim 18 is disclosed by Westbrook at element 415, Fig. 4A; Fig. 8; and at column 13, line 53 through column 14, line 12. (Office Action, page 6). Figs. 4A (element 415) and Fig. 8 of Westbrook illustrate a queue manager. These figures in no way disclose or suggest the reorder components recited in claim 18. Column 13, line 53 through column 14, line 12 of Westbrook states:

When all these sub-packets have been received by one or more of the distributed packet reassemblers, this information is communicated to

Serial No.: 10/024,553
Docket No.: 0023-0062

the particular packet reassembler holding the head of the packet. The data structure is then forwarded over link 411 to the corresponding queue manager, such as queue manager 415, to store the information in a queue corresponding to the destination of the reassembled packet. The operation of one embodiment of queue manager 415 is further described in relation to FIG. 8. Queue manager 415 receives the description of the reassembled packet, temporarily stores it in the incoming buffer 802, and then stores it in queue memory 806 in a queue based on its destination (and possibly priority and/or class of service). At the appropriate time, as determined by control logic 808, the queue manager extracts from one of its queues a data structure describing the corresponding reassembled packet to be send from the distributed resequencing and reassembly component 303B, and places it in outgoing buffer 804, which is then forwarded back to packet reassembler 410 over link 412.

Packet reassembler 410 receives a pointer to the data structure reflecting the reassembled packet from queue manager 415. The information in this data structure is forwarded to packet memory manager 420. Packets comprising the reassembled packet are placed on the packet merge bus 305B-E at the appropriate time to generated the reassembled packet out packet merge bus 305E.

(Westbrook, col. 13, line 32 through col. 14, line 12). This section of Westbrook describes the operation of queue manager 415 and packet reassembler 410 in reassembling a packet. This section of Westbrook, however, completely fails to disclose or suggest, as is recited in claim 18, classifying data items based on a sequence number of the data items and based on a position of the reorder buffer relative to the sequence number space, the reorder engine determining whether to store the data items in the reorder buffer based on the classification of the data items.

In interpreting the above-quoted section of Westbrook, the Examiner states that "[t]he reassembled packet is stored in queue memory 806 based on its destination, priority and/or class of service." (Office Action, page 6). Applicant agrees with the Examiner in so much that Westbrook discloses storing a data

Serial No.: 10/024,553
Docket No.: 0023-0062

structure "in a queue corresponding to the destination of the reassembled packet."

(Westbrook, col. 13, lines 58-60). Storing a packet in a queue based on its destination, however, is not reasonably related to, as is recited in claim 18, classifying a data item based on the sequence number of the data item and based on a position of the reorder buffer relative to the sequence number space.

Westbrook completely fails to disclose or suggest any such classification. Further, Westbrook clearly does not then determine whether to store the data items in the reorder buffer based on the classification of the data items, as is also recited in claim 18.

For at least these reasons, Applicant submits that Westbrook does not disclose or suggest each of the features recited in claim 18, and therefore, the rejection of this claim should be withdrawn. The rejection of claims 19-24 based on Westbrook should also be withdrawn, at least by virtue of the dependency of these claims from claim 18.

Claims 19-24 recite features of their own that are not disclosed or suggested by Westbrook.

For example, claim 23 recites that the classification utilizes a plurality of regions, including "a first region corresponding to the sequence number of the data item falling within a sequence number range covered by the reorder buffer, and a second region corresponding to the sequence number of the data item falling within a sequence number range extending from the value stored in the reorder buffer pointer to a predetermined range beyond the value stored in the reorder buffer pointer." Westbrook does not disclose or suggest classifying data items using the

Serial No.: 10/024,553
Docket No.: 0023-0062

first and second regions recited in claim 23. In rejecting claim 23, the Examiner points to different locations 443A and 443B within ring buffer 442 of Westbrook. Applicant submits that simply because Westbrook discloses the use of a ring buffer in no way discloses or suggests the features of claim 23 (and the claims from which it depends), which include classifying data units into the specific regions recited in claim 23. The mere disclosure of a ring buffer that includes a number of buckets in no way discloses or suggests classification into the regions specifically recited in claim 23. Accordingly, Westbrook also does not disclose or suggest the features of claim 23.

Independent claim 26 is directed to a system for ordering data items, including means for receiving the data items, the received data items being associated with a sequence number chosen from a sequence number space, the sequence numbers indicating a relative order of the data items. The system also includes means for classifying the received data items into one of a plurality of possible classifications based on the sequence number of the data items in relation to a value that defines the position of a reorder buffer in the sequence number space. Further, the system includes means for processing the data items in the reorder buffer based on the classifications of the data items.

Westbrook does not disclose or suggest each of the features recited in claim 26. Westbrook, for example, does not disclose or suggest "means for classifying the received data items into one of a plurality of possible classifications based on the sequence number of the data items in relation to a value that defines the position of a reorder buffer in the sequence number space." In rejecting this claim, the

Serial No.: 10/024,553
Docket No.: 0023-0062

Examiner points to the section of Westbrook that was discussed above with regard to the rejection of claim 18. (Office Action, page 9). This section of Westbrook describes the operation of queue manager 415 and packet reassembler 410 in reassembling a packet. This section of Westbrook, however, completely fails to disclose or suggest the means for classifying, as recited in claim 26.

Accordingly, Westbrook also fails to disclose or suggest each of the features recited in claim 26. Therefore, the rejection of claim 26 is also improper and should be withdrawn.

Rejection of Claim 25

Claim 25 stands rejected under 35 U.S.C. § 103(a) as being unpatentable over Westbrook in view of Kyker.

Applicant has reviewed Kyker and submits that Kyker does not cure the above-discussed deficiencies of Westbrook with regard to claim 18. Accordingly, the rejection of claim 25 is improper and should be withdrawn.

Conclusion

In view of the foregoing amendments and remarks, Applicant respectfully requests withdrawal of the outstanding rejections and the timely allowance of this application.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account 50-1070

Serial No.: 10/024,553
Docket No.: 0023-0062

and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY SNYDER, L.L.P.

By:



Brian Ledell
Reg. No. 42,784

Date:
11350 Random Hills Road
Suite 600
Fairfax, VA 22030
Telephone: (571) 432-0800
Facsimile: (571) 432-0808

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

☐ **BLACK BORDERS**

☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**

☐ **FADED TEXT OR DRAWING**

☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**

☐ **SKEWED/SLANTED IMAGES**

☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**

☐ **GRAY SCALE DOCUMENTS**

☒ **LINES OR MARKS ON ORIGINAL DOCUMENT**

☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**

☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.